

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:
Sompong Paul Olarig et al.

Serial No.: 10/039,010

Filed: December 31, 2001

For: Supporting Interleaved Read/Write
Operations from/to Multiple Target
Devices

§ Confirmation No.: 7506
§
§ Group Art Unit: 2188
§
§ Examiner: Chery, Mardochee
§
§
§ Atty. Docket: NUHP:0107
§ 200304299-1
§

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

CERTIFICATE OF TRANSMISSION OR MAILING
37 C.F.R. 1.8

I hereby certify that this correspondence is being transmitted by facsimile to the United States Patent and Trademark Office in accordance with 37 C.F.R. § 1.6(d), or is being transmitted via the Office electronic filing system in accordance with 37 C.F.R. § 1.6(a)(4), or is being deposited with the U.S. Postal Service with sufficient postage as First Class Mail in an envelope addressed to: Mail Stop Appeal Brief-Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date below:

September 22, 2008

Date

Jeffery R. Peterson

AMENDED APPEAL BRIEF
PURSUANT TO 37 C.F.R. §§ 41.31 AND 41.37

This Amended Appeal Brief is being filed in response to the Notification of Non-Compliant Appeal Brief mailed July 22, 2007 and replaces the Appeal Brief filed May 1, 2008. Appellants assert that this Amended Appeal Brief is complete and meets the requirements of 37 C.F.R. § 41.37(c).

The Commissioner is authorized to charge \$120.00, for a one-month extension, and any additional fees which may be necessary to advance prosecution of the present application, to Account No. 08-2025, Order No. 200304299-1.

1. **REAL PARTY IN INTEREST**

The real party in interest is Hewlett-Packard Development Company, L.P., the Assignee of the above-referenced application by virtue of the Assignment recorded at reel

13085, frame 0327, and dated July 8, 2002. Accordingly, Hewlett-Packard Development Company, L.P. will be directly affected by the Board's decision in the pending appeal.

2. **RELATED APPEALS AND INTERFERENCES**

Appellants are unaware of any other appeals or interferences related to this Appeal. The undersigned is Appellants' legal representative in this Appeal.

3. **STATUS OF CLAIMS**

Claims 1-6, 8-11, 13, 16-18, 20-27, 29, and 32-48 are currently pending and under rejection. Claims 1-6, 8-11, 13, 16-18, 20-27, 29, 32-48 are the subject of this Appeal. Claims 7, 12, 14, 15, 19, 28, 30, and 31 were previously canceled.

4. **STATUS OF AMENDMENTS**

As the instant claims have not been amended since the mailing of the Office Action, there are no outstanding amendments to be considered by the Board.

5. **SUMMARY OF CLAIMED SUBJECT MATTER**

The present invention relates generally to read/write transactions on a computer bus and, more particularly, to a method and apparatus for supporting interleaved read/write operations for multiple target devices in a multicast computer environment. The Application contains eight independent claims, namely, claims 1, 9, 16, 25, 32, 33, 36, and 47, all of which are the subject of this Appeal. The subject matter of these claims is summarized below.

With regard to the aspect of the invention set forth in independent claim 1, discussions of the recited features of claim 1 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 1 generally recites a method for transacting between an initiator device and a plurality of target devices. *See, e.g.,* specification, page 8, lines 5-6; Figs. 4 and 5. The method includes associating

each of the plurality of target devices with a single base address, wherein the same single base address is associated with each of the plurality of target devices. *Id.* at page 8, lines 6-9, and 21-23; Figs. 3A, 4, and 5. The method also includes sending a multicast transaction from the initiator device to the plurality of target devices, wherein sending the multicast transaction comprises sending a multicast transaction to the single base address associated with each of the plurality of target devices. *Id.* at page 8, lines 5-9; page 11, lines 2-4 and 31-34; Figs. 4 and 5. Additionally, the method includes executing the transaction when the transaction is received by the plurality of target devices according to the configuration of the target devices. *Id.* at page 10, lines 7-22; page 11, line 31 through page 12, line 13; Figs. 4 and 5.

With regard to the aspect of the invention set forth in independent claim 9, discussions of the recited features of claim 9 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 9 generally recites a method for transacting data stored in memory between an initiator device and multiple target devices. *See, e.g.*, specification, page 8, lines 5-6; Figs. 4 and 5. The method includes detecting a multicast transaction request. *Id.* at page 8, lines 5-9 and 21-23; page 10, lines 9-12; page 11, lines 2-4 and 31-34; Figs. 4 and 5. The method also includes accessing a first portion of memory by a first target device in response to the multicast transaction request and accessing a second portion of memory by a second target device concurrently with the access to the first portion of memory in response to the multicast transaction request. *Id.* at page 9, lines 22-32; page 10, lines 13-22; page 12, lines 1-13; Figs. 2, 3B-C, 4, and 5. The first and second portions of memory are accessed with a single base address associated with both the first target device and the second target device, wherein the first target device and the second target device are associated with the same single base address. *Id.* at page 8, lines 5-9 and 21-23; page 9, lines 9-12; page 10, lines 9-11; Figs. 4 and 5.

With regard to the aspect of the invention set forth in independent claim 16, discussions of the recited features of claim 16 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 16 generally recites a computer system including a bus. *See, e.g.*, specification, page 5, lines 20-33; Fig. 1. An initiator device is coupled to the bus, the initiator device configured to initiate a transaction request. *Id.* at page 6, lines 1-4 and 19-26; page 7, lines 27-28; page 8, lines 5-6; Fig. 1. A plurality of target devices coupled to the bus, wherein each of the plurality of target devices concurrently executes a portion of the transaction request. *Id.* at page 6, lines 6-9 and 20-23; page 8, lines 5-11; page 9, line 22 through page 10, lines 5 and 24-25; page 12, lines 9-13 and 22-25; Fig. 1. Additionally, the initiator device is configured to multicast the transaction request to the plurality of target devices using a single base address associated with the plurality of target devices, wherein the same single base address is associated with each of the plurality of target devices. *Id.* at page 8, lines 5-9 and 21-23; page 9, lines 9-12; page 10, lines 9-11; Figs. 1, 4 and 5.

With regard to the aspect of the invention set forth in independent claim 25, discussions of the recited features of claim 25 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 25 generally recites a computer system including: a processor and a bus coupled to the processor. *See, e.g.*, specification, page 5, lines 22-33; Fig. 1. An initiator device is coupled to the bus, the initiator device is configured to issue a multicast transaction. *Id.* at page 5, lines 22-25; page 6, lines 19-20; page 8, lines 5-6; Fig. 1. A plurality of target devices are coupled to the bus, the plurality of target devices are configured to execute the multicast transaction with concurrent data responses from a plurality of interleaved memory regions. *Id.* at page 5, lines 24-33; page 9, line 22 through page 10, line 5 and lines 9-12 and 18-28; Fig. 1 and 2B. The initiator device is configured to multicast the transaction request to the plurality of target devices using a single base address associated with the plurality of target devices, wherein the same single base address is associated with each

of the plurality of target devices. *Id.* at page 8, lines 5-9 and 21-23; page 9, lines 9-12; page 10, lines 9-11; Fig. 1.

With regard to the aspect of the invention set forth in independent claim 32, discussions of the recited features of claim 32 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 32 generally recites a computer including a memory and a controller configured to logically divide the memory into a plurality of interleaved memory regions. *See, e.g.*, specification, page 5, lines 20-24; page 6, lines 1-4; page 9, lines 18-32; Figs. 1 and 3A-C. The computer has a plurality of devices having a common base address, wherein each of the plurality of devices is associated with one of the interleaved memory regions and wherein each of the devices simultaneously accesses its associated interleaved memory region in response to a single transaction request directed to the common base address. *Id.* at page 8, lines 5-11 and 21-23; page 9, lines 22 through page 10, line 16; Figs. 1, 2B and 3A-C.

With regard to the aspect of the invention set forth in independent claim 33, discussions of the recited features of claim 33 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 33 generally recites a method including dividing a section of memory into a plurality of interleaved memory regions and associating the plurality of interleaved memory regions with a plurality of target devices. *See, e.g.*, specification, page 9, line 9 through page 10, line 5; Figs. 1 and 3A-C. The method also includes associating the plurality of target devices with a single base memory address, wherein the same single base memory address is associated with each of the plurality of target devices and executing a memory access using the single base memory address. *Id.* at page 8, lines 21-31; page 9, lines 9-12; page 10, lines 11-32; Figs. 1, 3A-C, 4 and 5.

With regard to the aspect of the invention set forth in independent claim 36, discussions of the recited features of claim 36 can be found at least in the below cited

locations of the specification and drawings. By way of example, claim 36 generally recites a tangible machine readable medium including code to initialize a plurality of devices and code to configure the plurality of devices to associate a single base address with the plurality of devices, wherein the same single base address is associated with each of the plurality of devices. *See, e.g.*, specification, page 8, lines 13-23; page 9, lines 9-12; Figs. 3A-C. Additionally, the tangible machine readable medium includes code to associate the single base address with a plurality of interleaved memory regions, wherein the same single base address is associated with each of the plurality of interleaved memory regions. *Id.* at page 9, lines 9 through page 10, line 32; Figs. 2B and 3A-C.

With regard to the aspect of the invention set forth in independent claim 47, discussions of the recited features of claim 47 can be found at least in the below cited locations of the specification and drawings. By way of example, claim 47 generally recites a method for performing operations accessing memory including assigning a common base address to a plurality of disk drives to create a collective target group. *See, e.g.*, specification, page 8, lines 5-31; page 9, lines 18-32; Figs. 2B and 3A-C. A portion of data storage is assigned to each disk drive of the collective target group and an interleaved memory operation is initiated as a single request to the collective target group using the common base address, wherein each disk drive recognizes the common base address and simultaneously executes the memory operation on the portion of data storage to which the disk drive is assigned. *Id.* at page 9, line 9 through page 10, line 32; Figs. 2B, 3A-C, 4 and 5.

A benefit of the invention, as recited in these claims, is the ability to save the resources of both an initiator device and a bus by having the target devices simultaneously execute requests and return data. *See* specification, page 8, lines 5-11. In particular, there is improved performance by avoiding wait states and limiting inefficiencies shifting data transfer operations from the executing device to the data bus.

Id. at page 6, lines 19-23. This is a clear difference and distinction from the prior art, as discussed below.

6. **GROUND OF REJECTION TO BE REVIEWED ON APPEAL**

First Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's first ground of rejection in which the Examiner rejected claim 32, 47 and 48 under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement.

Second Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's second ground of rejection in which the Examiner rejected claims 1-6, 8-11, 13, 16-18, 20-23, 25-27, and 29 under 35 U.S.C. § 103(a) as being unpatentable over Leung et al., U.S. Patent No. 6,272,577 (hereinafter "Leung") in view of Guttag, U.S. Patent No. 5,761,726 (hereinafter "Guttag").

Third Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's third ground of rejection in which the Examiner rejected claim 24 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Carmichael, et al., U.S. Patent No. 5,864,712 (hereinafter "Carmichael").

Fourth Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's fourth ground of rejection in which the Examiner rejected claims 32 and 47 under 35 U.S.C. § 103(a) as being unpatentable over Guttag and Gupta, U.S. Patent No. 6,405,286 (hereinafter "Gupta") in view of Blaner, U.S. Patent No. 5,737,575 (hereinafter "Blaner").

Fifth Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's fifth ground of rejection in which the Examiner rejected claim 48 under 35 U.S.C. § 103(a) as being unpatentable over Guttag and Gupta in view of Blaner and further in view of Olarig, U.S. Patent No. 6,230,225 (hereinafter "Olarig").

Sixth Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's sixth ground of rejection in which the Examiner rejected claims 33-38 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Gupta.

Seventh Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's seventh ground of rejection in which the Examiner rejected claims 39-44 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Olarig.

Eighth Ground of Rejection for Review on Appeal:

Appellants respectfully urge the Board to review and reverse the Examiner's eighth ground of rejection in which the Examiner rejected claims 45 and 46 under 35 U.S.C. § 103(a) as being unpatentable over Gupta in view of Blaner and further in view of Olarig.

7. **ARGUMENT**

As discussed in detail below, the Examiner has improperly rejected the pending claims. Further, the Examiner has misapplied long-standing and binding legal precedents and principles in rejecting the claims under Section 103. Accordingly, Appellants respectfully request full and favorable consideration by the Board, as Appellants strongly

believe that claims 1-6, 8-11, 13, 16-18, 20-27, 29, and 32-48 are currently in condition for allowance.

A. **Ground of Rejection No. 1:**

The Examiner rejected claims 32, 47, and 48 under 35 U.S.C. § 112, first paragraph. Specifically, the Examiner stated:

Claims 32 and 47-48 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s) had possession of the claimed invention. In particular, claim 32 recites inter alia a plurality of devices “having a common base address” and a single transaction request “directed to the common base address”. Claims 47 recites inter alia “assigning a common base address to a plurality of disk drives to create a collective target group; initiating an interleaved memory operation . . . using the common base address, wherein each disk drive recognizes the common base address and simultaneously executes the memory operation on the portion of data storage to which the disk drive is assigned”. Although the original disclosure provides for “a single base address”, it does not provide for these newly added limitations.

Office Action, page 6.

Legal Precedent

First, regarding the written description requirement, the initial burden of proof regarding the sufficiency of the written description falls on the Examiner. Accordingly, the Examiner must present evidence or reasons why persons skilled in the art would not recognize a description of the claimed subject matter in the applicant’s disclosure. *In re Wertheim*, 541 F.2d 257, 262, 191 U.S.P.Q. 90, 96 (C.C.P.A. 1976). Additionally, the written description requirement does not require the claims to recite the same terminology used in the disclosure. The patentee may be his own lexicographer. *Ellipse*

Corp. v. Ford Motor Co., 171 U.S.P.Q. 513 (7th Cir. 1971), *aff'd*, 613 F.2d 775 (7th Cir. 1979), *cert. denied*, 446 U.S. 939 (1980). Furthermore, breadth of a claim should not be equated with indefiniteness. *In re Miller*, 441 F.2d 689, 169 U.S.P.Q. 597 (CCPA 1971).

Second, the pending claims must be given an interpretation that is reasonable and consistent with the *specification*. See *In re Prater*, 415 F.2d 1393, 1404-05, 162 U.S.P.Q. 541, 550-51 (C.C.P.A. 1969) (emphasis added); see also *In re Morris*, 127 F.3d 1048, 1054-55, 44 U.S.P.Q.2d 1023, 1027-28 (Fed. Cir. 1997); see also M.P.E.P. §§ 608.01(o) and 2111. Indeed, the specification is “the primary basis for construing the claims.” See *Phillips v. AWH Corp.*, No. 03-1269, -1286, at 13-16 (Fed. Cir. July 12, 2005) (*en banc*). Furthermore, interpretation of the claims must also be consistent with the interpretation that *one of ordinary skill in the art* would reach. See *In re Cortright*, 165 F.3d 1353, 1359, 49 U.S.P.Q.2d 1464, 1468 (Fed. Cir. 1999); M.P.E.P. § 2111. “The inquiry into how a person of ordinary skill in the art understands a claim term provides an objective baseline from which to begin claim interpretation.” See *Collegenet, Inc. v. ApplyYourself, Inc.*, 418 F.3d 1225, 75 U.S.P.Q.2d 1733, 1738 (Fed. Cir. 2005) (quoting *Phillips v. AWH Corp.*, 75 U.S.P.Q.2d 1321, 1326).

Claim 32 complies with the first paragraph of Section 112.

Appellants respectfully assert that claims 32, 47 and 48 fully comply with the written description requirement of the first paragraph of Section 112. In particular, the specification provides support for the claimed feature a “common base address,” as set forth in claims 32, 47 and 48. The specification states, “target groups are collectively configured with a *single base address*. This allows a grouping of targets to recognize an initiator request with a *single base address* as a request to all of the target devices 150 included in the target group.” Specification, page 8, lines 21-23 (emphasis added). As such, multiple targets (“grouping of targets”) recognize a *single base address* request because the single base address is common or shared with each target in the grouping of targets. See *id.*, e.g., at page 9, lines 9-10; page 10, lines 9-11; page 11, lines 8-10.

Moreover, Appellants respectfully assert that this would be readily understood to those of ordinary skill in the art, despite the lack of the specific term “common” in the specification. Because the single base address is recognized and shared by the target devices, it may be considered “common” to each of the target devices. Accordingly, the specification clearly sets forth that each of the target devices of a group shares a “common base address,” as set forth in the claims, contrary to the Examiner’s assertion and the rejection under Section 112, first paragraph, as failing to comply with the written description requirement is wholly without merit. Appellants, therefore respectfully request reversal of the rejection of claim 32, 47 and 48 under Section 112.

Furthermore, it appears that the Examiner is suggesting that the recitation of “common base address” adds new matter. However, in view of the foregoing discussion, it is clear that there is more than adequate support set forth in the specification for the use of the terms and that the use of the terms does not constitute new matter. Indeed, the Examiner’s assertion that the use of the term “common” adds new matter may indicate that the Examiner is not interpreting the claims in view of the specification and as one of ordinary skill in the art would interpret the claims, as required by law. In particular, Appellants respectfully assert that one of ordinary skill in the art would consider the use of the term “common” germane within the context that it is used in the claims in view of the specification. As such, for at least this additional reason, Appellants respectfully request reversal of the rejection under Section 112 of claim 32, 47 and 48.

B. Ground of Rejection No. 2:

The Examiner rejected claims 1-6, 8-11, 13, 16-18, 20-23, 25-27, and 29 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag. Specifically, with respect to claims 1, 9, 16 and 25, the Examiner stated:

As per claim 1, Leung et al. discloses a method for transacting between an initiator device and a plurality of target devices [*a memory device in which a single input data stream can be simultaneously written to multiple*

memory arrays; col. 3, lines 63-65]; sending a multicast transaction from the initiator device to the plurality of target devices, wherein sending the multicast transaction comprises sending a multicast transaction to the single base address associated with each of the plurality of target devices [a base address which identifies the memory module; col. 10, lines 20-23; the memory modules are equipped with independent address and command decoders so that they function as independent units, each with their own base address; the present invention groups at least two memory arrays or banks into a memory module and connects all the memory modules to a bus; col. 4, lines 20-24; in each memory module, a programmable identification register contains the base address of the memory module and a mechanism which decommisions the module from acting on certain memory access; col. 5, lines 5-8; col. 4, lines 31-33; multiple bank operations such as broadcast-write and interleaved access possible; a memory device able to handle a broadcast write operations simultaneously; col. 5, lines 27-31]; executing the transaction when the transaction is received by the plurality of target devices according to the configuration of the target devices [since each memory module is a complete functional unit, the memory module architectures allows parallel processes and multiple memory module operations to be performed within different memory modules; col. 4, lines 42-45].

However, Leung does not explicitly teach associating each of the plurality of target devices with a single base address, wherein the same single base address is associated with each of the plurality of target devices as recited in the claim.

Guttag discloses associating each of the plurality of target devices with a single base address, wherein the same single base address is associated with each of the plurality of target devices [col. 172, lines 48-55] to generate addresses for read/write access to data stored within a plurality of memories (col. 5, ll. 40-45).

Since the technology for implementing a data processing system with associating each of the plurality of target devices with a single base address, wherein the same single base address is associated with each of the plurality

of target devices was well known as evidenced by Guttag, an artisan would have been motivated to implement this feature in the system of Leung in order to generate addresses for read/write access to data stored within a plurality of memories. Thus it would have been obvious to one of ordinary skill in the art at the time of invention by Applicant to modify the system of Leung to include associating each of the plurality of target devices with a single base address, wherein the same single base address is associated with each of the plurality of target devices because this would have generated addresses for read/write access to data stored within a plurality of memories (col. 5, ll 40-45) as taught by Guttag.

As per claim 9, the rationale in the rejection of claim 1 is herein incorporated. Leung et al. further discloses, a method for transacting data stored in memory between an initiator device and detecting a multicast transaction request [*multiple bank operations such as broadcast-write and interleaved access possible; a memory device able to handle a broadcast write bandwidth of 36 gigabytes per second and 36 memory operations simultaneously; col. 5, lines 27-31*] accessing a first portion of memory by a first target device associated with the first portion of memory in response to the muticast transaction request [*when a memory read or write command is decoded, each memory module examines the communication ID of the command. All modules, except the module to which the command is addressed, go into an idle state until the read or write transaction is finished; col. 19, lines 42-47*]; accessing a second portion of memory by a second target device associated with the second portion of memory concurrently with access to the first portion of memory in response to the multicast transaction request wherein the first and second portions of memory are accessed with a single base address associated with the first target device and the second target device [*a first field contains a base address which identifies the memory module by communication address. A second field contains an address which identifies the memory array within the memory module. Col. 10, lines 21-25; the present invention groups at least two memory arrays or banks into a memory module and connects all the memory modules to a bus; col. 4, lines 20-24; in each memory module, a programmable*

identification register contains the base address of the memory module and a mechanism which decommissions the module from acting on certain memory access; col. 5, lines 5-8]

Office Action, pages 7-9.

Legal Precedent

Appellants respectfully traverse this rejection. The burden of establishing a *prima facie* case of obviousness falls on the Examiner. *Ex parte Wolters and Kuypers*, 214 U.S.P.Q. 735 (B.P.A.I. 1979). To establish *prima facie* obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art. *In re Royka*, 180 U.S.P.Q. 580 (C.C.P.A. 1974). However, it is not enough to show that all the elements exist in the prior art since a claimed invention composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art. *KSR International Co. v. Teleflex Inc.*, 127 S.Ct. 1727, 1741 (2007). It is important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does. *Id.* Specifically, there must be some articulated reasoning with a rational underpinning to support a conclusion of obviousness; a conclusory statement will not suffice. *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006). Indeed, the factual inquiry determining whether to combine references must be thorough and searching, and it must be based on *objective evidence of record*. *In re Lee*, 61 U.S.P.Q.2d 1430, 1436 (Fed. Cir. 2002).

Independent Claim 1

Independent claim 1 recites, *inter alia*, “A method for transacting between an initiator device and a plurality of target devices, the method comprising: associating each of the plurality of target devices with *a single base address*, wherein *the same single base address is associated with each of the plurality of target devices*; sending a *multicast transaction from the initiator device to the plurality of target devices*, wherein sending

the multicast transaction comprises sending a multicast transaction to *the single base address* associated with each of the plurality of target devices; sending a multicast transaction from the initiator device to the plurality of target devices, wherein sending the multicast transaction comprises sending a multicast transaction to the single base address associated with each of the plurality of target devices.” (Emphasis added).

In sharp contrast to the claimed subject matter, the Leung and Gutttag references, taken alone or in combination, fail to disclose several features of recited in claim 1. In particular, the Leung and Gutttag references fail to disclose sending a *multicast transaction* from an *initiator device* to a plurality of *target devices*. Additionally, neither reference discloses the same single address being associated with each of the plurality of targets. These deficiencies will be discuss in detail below.

As discussed in the application, in accordance with the teachings of the present application, an initiator, i.e., CPU 110, may issue a single request seen by multiple targets, i.e., PCI devices 150 or controllers, of a target grouping. *See, e.g.*, application, pg 6, lines 6-17; pg. 8, lines 5-11; Figs. 1, 2A, 2B, 4A and 4B. To accomplish this, each target device of a target group is configured with a single base address. *See id.* at page 8, lines 21-23. The entire target group then listens to the request and each target processes only its portion of the request simultaneous with each of the other targets. *See id.* pg. 10, lines 11-16; Fig. 4.

In applying the Leung and Gutttag references to reject claim 1, the Examiner has mistakenly misconstrued the claim, misinterpreted the teachings of the cited references and, furthermore, been inconsistent in applying the references. In particular, in rejecting claim 1, the Examiner asserts that a “memory module” of Leung is the equivalent of the “target device” of claim 1. While it is true that Leung discloses a plurality of memory modules, in view of the specification, one of ordinary skill in the art would not construe the target devices recited in claim 1 as being a memory module as disclosed by Leung. In

particular, as mentioned above, the instant application discloses the target device to be PCI devices or controllers. *See* application, pg. 6, lines 6-17, Fig. 1. The memory modules of Leung, in contrast and as characterized by the Examiner, are at least two memory banks or arrays equipped with independent address and command decoders. *See* Office Action, pg. 7. As such, one of ordinary skill in the art would not understand the memory modules set forth in Leung to be the same as the target devices of claim 1.

Additionally, while Appellants recognize that the claims are given the broadest reasonable interpretation during examination, Appellants assert that the interpretation given by the Examiner is simply not reasonable. The unreasonableness of the Examiner's interpretation is evident in the inconsistency of the Examiner's treatment of the cited references. Specifically, while the Examiner asserts that the memory modules of Leung are equivalent to the target devices of claim 1, in rejecting claims 5 and 6, the Examiner cited to an input/output controller and a memory controller, respectively, of Leung as being the target devices. *See* Office Action, pg. 11. As such, the Examiner has inconsistently identified both memory modules and controllers as being the "plurality of target devices." This clearly is not reasonable, as they serve different purposes and cannot reasonably be considered to be the same thing, i.e., a controller is not a memory module and vice-versa. Moreover, even if, *arguendo*, the input/output controller and/or the memory controller of Leung could reasonably be considered to the target device, Leung only discloses a single target device. As such, even if the controller of Leung were considered to be a "target device," Leung would still fail to disclose a "*plurality* of target devices," as set forth in claim 1.

Guttag suffers from the same deficiencies as Leung in this regard. Specifically, Guttag discloses a plurality of processors and a plurality of memories. *See* Guttag, col. 5, lines 35-36. However, there is nothing that can reasonably be considered the same as a plurality of target device as set forth in the instant application and recited in claim 1. Under the same rationale that the memory devices of Leung cannot reasonably be

considered target devices, the memories of Guttag cannot be considered target devices. As such, Guttag fails to overcome the deficiencies of Leung with respect to this element of claim 1.

Furthermore, with respect to the same single address being associated with each of the plurality of targets, Appellants agree with the Examiner that Leung fails to disclose the feature of the plurality of target devices having a single base address, wherein the *same single base address is associated with each of the plurality of target devices*. However, contrary to the Examiner's assertion, Guttag does not overcome this admitted deficiency of Leung. In particular, Guttag does not disclose a plurality of target devices with "*a single base address, wherein the same single base address is associated with each of the plurality of target devices,*" as set forth in claim 1. Rather, Guttag discloses unique addresses, i.e., different address for each of a plurality of memories. As stated in Guttag:

Each of the memories has a unique addressable memory portion of a single memory space. Each processor has a predetermined plurality of corresponding memories. These corresponding memories have a corresponding base address within said single memory address space.

Guttag, at col. 5, lines 36-41 (emphasis added). In rejecting claim 1, the Examiner referred to col. 172, lines 48-55 (claim 1) of Guttag as disclosing the above-recited claim features. In its entirety, the cited portion states:

a plurality of n processors, where n is less than m and each of said n processors has a predetermined *plurality of corresponding memories*, said predetermined plurality of memories corresponding to *each processor having a corresponding fixed base address* within said single memory address space, each of said processors capable of generating any address within said single memory address space for read/write access to data stored within said plurality of m memories.

Guttag, col. 172, lines 48-55 (emphasis added). As the above-quoted sections clearly illustrate, Guttag discloses a plurality of processors that each have *their own* (i.e., *corresponding*) base address within *their own individual section* of the memory address space. *See also*, Guttag, col. 172, lines 43-47. That is, a memory space is divided into a plurality of *m* memories, each memory having a “*unique* addressable memory portion”. Consequently, the claim feature “wherein *the same single base address* is associated with each of the plurality of target devices” is not taught by Guttag. Indeed, Guttag teaches the antithesis of this feature, in that each memory has its own, *unique* base address.

Additionally, Appellants respectfully assert that the Leung and Guttag references do not disclose anything with respect to *multicast transactions*, as set forth in claim 1. In particular, the Leung and Guttag references disclose nothing with respect to “sending a *multicast transaction* from the initiator device to the plurality of target devices, wherein sending the *multicast transaction* comprises sending a *multicast transaction to the single base address associated with each of the plurality of target devices*,” as recited in claim 1. With respect to this feature, the Examiner has only cited to a single line in Leung which indicates the possibility of broadcast-write. Leung, however, explains that the broadcast-write operation is conducted by setting an array-select bit of a memory device. *See* Leung col. 23, line 52 through col. 24, line 8. Each memory device that has the array-select bit set participates in the write operation. *Id.* The select bit is not an address. Moreover, one of ordinary skill in the art would not consider the select bit to be an address, much less a *single base address associated with each of the plurality of target devices*, as set forth in claim 1. As such, the broadcast-write of Leung cannot reasonably be considered as disclosing the multicast transaction as set forth in claim 1. Guttag fails to overcome this deficiency.

In view of these clear deficiencies, the Leung and Guttag references, taken alone or in combination, fail to disclose all the features of claim 1. Appellants respectfully assert, therefore, that the Examiner has not established a *prima facie* case of obviousness

with regard to claim 1. Accordingly, Appellants respectfully request reversal of the Section 103 rejection and allowance of independent claim 1 and the claims that depend therefrom.

Dependent claim 2

Claim 2 recites, “assigning a base memory address to be shared by the plurality of target devices; and assigning a first portion of memory to a first target device of the plurality of target devices.” In rejecting claim 2, the Examiner cites to Leung. However, in sharp contrast to claim 2, Leung discloses that each of the memory modules is independent and has its *own base address*. See Leung, col. 4, lines 32-35. As such, Leung fails to disclose “assigning a base memory address to be *shared* by the plurality of target devices,” as set forth in claim 2. Accordingly, for at least this reason, Appellants assert that claim 2 is not rendered obvious in view of Leung taken alone or in combination with Gutttag. Appellants respectfully request reversal of the rejection of claim 2.

Dependent claim 3

Dependent claim 3 recites, *inter alia*, “initiating a read operation by the plurality of target devices assigned to the base memory address; fetching stored data from a portion of memory associated with each of the target devices, *the data being concurrently fetched by each associated target device*; and sending the fetched data to the initiator device.” Dependent claim 4 recites, *inter alia*, “initiating a write operation by the plurality of target devices assigned to the base memory address; and writing data of the write request to a portion of memory associated with each target device, *the data being concurrently written by each associated target device*.”

In rejecting claims 3 and 4, the Examiner cited to portions of Leung that disclose reading or writing to different arrays in a *time multiplexed* manner. See Office Action, pg. 10. Indeed, Leung only discloses reading and/or writing to the memory modules in a

time multiplexed manner. *See* Leung, col. 24, lines 23-27. Time multiplexing indicates that they events occur at different moments of time, i.e., not concurrent. Accordingly, the subject matter of claims 3 and 4 is not made obvious under Section 103 in view of Leung when taken alone or in combination with Guttag. Accordingly, Appellants request reversal of the rejection of claims 3 and 4.

Dependent claims 5 and 6

Claim 5 sets forth wherein the target devices comprise “input/output controllers.” Claim 6 sets forth wherein the target devices comprise “disk array controllers.” As discussed above, in rejecting claim 1, the Examiner cited to the memory modules of Leung as being target devices. In rejecting claims 5 and 6, the Examiner cited to the I/O module 104 and the controller 1920 as being the target devices. As such, the Examiner is inconsistently treating the references as well as the claims. This is improper and does not reflect an interpretation that one of ordinary skill in the art would give. In particular, one of ordinary skill in the art would consistently interpret the references and the claims. For at least this reason, the Appellants respectfully request reversal of the rejection of claim 5 and 6.

Moreover, however, even assuming *arguendo* that the I/O module 104 or the controller 1920 could reasonably be considered “target devices,” Leung fails to disclose a plurality of I/O modules 104 or a plurality of controller 1920 and, as such, Leung would still be deficient in disclosing all the elements of claims 5 and 6. Accordingly, Appellants respectfully request reversal of the rejection of claim 5 and 6.

Dependent claim 8

Dependent claim 8 recites, “The method of claim 1, comprising a plurality of target groups.” As discussed above, the Examiner has cited to the memory modules of Leung as being the same as the target devices of claim 1. In rejecting claim 8, the Examiner cites to the two memory arrays of each memory module as being the equivalent

of the plurality of target groups of claim 8. However, the Examiner's interpretation is simply incorrect.

In particular, as discussed in the specification of the instant application, a target grouping includes target devices that have a single base address. *See* specification, paragraphs 26-28. Even assuming *arguendo* that the Examiner's interpretation of the memory modules as being target devices is accurate, the arrays of the memory modules cannot reasonably be considered to be a "plurality of target groups." Indeed, the arrays are simply parts of a single target device per the Examiner's reading, as such the two arrays taken together simply produce a single target device, not a plurality of target groups, as set forth in claim 8. Accordingly, for at least this reason, Appellants respectfully request reversal of the rejection of claim 8.

Independent claim 9

Independent claim 9 recites, *inter alia*, "*accessing a first portion of memory by a first target device in response to the multicast transaction request; and accessing a second portion of memory by a second target device concurrently with the access to the first portion of memory in response to the multicast transaction request, wherein the first and second portions of memory are accessed with a single base address associated with both the first target device and the second target device, wherein the first target device and the second target device are associated with the same single base address.*" (Emphasis added).

In sharp contrast, the Leung and Gutttag references, taken alone or in combination, fail to disclose several features of recited in claim 9. In particular, the Leung and Gutttag references fail to disclose "*accessing a first portion of memory by a first target device; and accessing a second portion of memory by a second target device concurrently.*" Additionally, neither reference discloses the first and second portions of memory are accessed with *a single base address associated with both the first target device and the*

second target device, wherein the first target device and the second target device are associated with the same single base address. These deficiencies are discussed in detail below.

As a preliminary matter, the Examiner has incorporated the rationale of the rejection of claim 1, but clearly only cited to Leung for the rejection of claim 9. As such, there is at least some ambiguity with respect to the rejection of claim 9 as to how Gutttag is being applied. Appellants request that the Examiner clarify his position on this rejection, as it appears to be incomplete. Otherwise Appellants respectfully request withdrawal of the rejection and allowance of the claim.

In rejecting claim 9, the Examiner again sets forth the “memory modules” of Leung as the equivalent of the “target devices” of claim 9. As discussed above, the memory modules cannot reasonably be considered the same as the target devices. In particular, as mentioned above, the instant application discloses the target device to be PCI device or controllers. *See* application, pg. 6, lines 6-17, Fig. 1. The memory modules of Leung, in contrast, are simply at least two memory banks or arrays equipped with independent address and command decoders. *See* Office Action, pg. 7. As such, one of ordinary skill in the art would not understand memory devices set forth in Leung and as characterized by the Examiner to be the same as the target devices of claim 1.

Additionally, the Examiner is inconsistent in the treatment of the claims and Leung. Specifically, in rejecting claim 1, the Examiner admitted that Leung does not disclose a plurality of target devices with a single base address. However, in rejection claim 9, the Examiner asserted that Leung discloses a single base address associated with a first and second target device by providing first and second address fields. *See* Office Action, pg. 20. Beyond this obvious inconsistency between the rejection of claim 1 and claim 9, this assertion is incorrect for several reasons. First, as set forth above, the memory modules are *not* target devices, as discussed above. Second, even assuming

arguendo that the memory modules could reasonably be considered target devices, the first and second field to which the examiner refers are pertinent only to accessing a particular array or bank within a *single* memory module. *See, e.g.*, Leung, col 10, lines 21-25. Indeed, each of the memory modules are independent and have their own base address. *Id.* at col. 4, lines 20-24. There is simply nothing in Leung that can reasonably be construed as disclosing a single base address associated with two memory modules, much less a single base address associated with a first target device and a second target device, as set forth in claim 9.

Moreover, another inconsistency of the Examiner is found with respect to claims 10 and 11. Specifically, while the Examiner asserts that the target devices are represented in Leung by the memory modules in rejecting claim 9, the Examiner cited to an input/output controller and a memory controller, respectively, of Leung as being the target devices in rejecting claims 10 and 11. *See* Office Action, pg. 12. As such, the Examiner has inconsistently identified both memory modules and controllers as being the “target devices.” This clearly is not reasonable, as they serve different purposes and cannot reasonably be considered to be the same thing in view of the disclosure of Leung, i.e., a controller is not a memory module and vice-versa. Furthermore, even if, *arguendo*, the input/output controller and/or the memory controller of Leung could reasonably be considered the target device, Leung only discloses a single controller, i.e., target device, and as such, fails to disclose a “second target device,” as set forth in claim 9.

In view of these clear deficiencies, Leung, taken alone or in combination with Guttag, fails to disclose all the features of claim 1. Appellants respectfully assert, therefore, that the Examiner has not established a *prima facie* case of obviousness with regard to claim 9. Accordingly, Appellants respectfully request reversal of the Section 103 rejection and allowance of independent claim 9 and the claims that depend therefrom.

Dependent claims 10 and 11

Claim 10 recites wherein the target devices comprise input/output controllers.

Claim 11 recites wherein the target devices comprise disk array controllers.

As discussed above, the Examiner has inconsistently treated the claim term “target device.” In particular, in rejecting claim 9, the Examiner indicated that the memory modules were the equivalent of the target devices. However, in rejecting claims 10 and 11, the Examiner equated an I/O module 104 and a controller 1920, respectively, as being target devices. As such, the Examiner has indicated both controllers and memory modules as being target devices. This is inconsistent and unreasonable. Indeed, one of ordinary skill in the art would not understand both the memory modules and controllers to be the same as target devices. Moreover, however, even assuming *arguendo* that the I/O module 104 or the controller 1920 could reasonably be considered “target devices,” Leung fails to disclose a plurality of I/O modules 104 or a plurality of controller 1920 and, as such, Leung would still be deficient in disclosing all the elements of claims 10 and 11. Accordingly, Appellants respectfully request reversal of the rejection of claim 10 and 11.

Dependent claim 13

Claim 13 recites, “accessing a plurality of target device, wherein the plurality of target devices are divided into a plurality of groups, wherein each of the plurality of groups is associated with a single base memory address configured to address the target devices within that group.”

Again, the Examiner has inconsistently applied Leung. In particular, in rejecting claim 13, it appears that the Examiner is intending that the arrays in each memory module to be the target devices, whereas in rejecting claim 9 the memory modules themselves were indicated as being the target devices. *See* Office Action, pg. 12. Such treatment is unreasonable, not consistent with what one of ordinary skill in the art would do, and

ignores the plain language of the claim. Specifically, if the memory modules are taken, *arguendo*, to be the target devices, then the memory modules then the claim calls for the memory modules to be divided into a plurality of groups, not portions of a single memory module, as suggested by the Examiner. Moreover, even under the Examiner's misguided and constrained interpretation, Leung fails to disclose wherein each of the plurality of groups is associated with a *single base memory* address, as set forth in claim 13. In particular, as noted by the Examiner, Leung discloses each of the memory modules has a unique base address and each address includes a first field that address a memory module (base address) and a second field that addresses an array in the module. *See* Leung, col. 4, lines 32-35; col. 10, lines 21-25. As such, each module has its own base address and each array has its own address, i.e., the second field. There is no reasonable reading of Leung that arrives at the subject matter covered by claim 13. As such, Appellants respectfully assert that the Examiner has not set forth a *prima facie* case for obviousness under Section 103. Accordingly, Appellants request reversal of the rejection of claim 13.

Independent claims 16 and 25

Independent claims 16 and 25 recite, *inter alia*, “an initiator device coupled to the bus, the initiator device configured to initiate a transaction request; and a plurality of target devices coupled to the bus wherein each of the plurality of target devices concurrently executes a portion of the transaction request, wherein the initiator device is configured to multicast the transaction request to the plurality of target devices using *a single base address associated with the plurality of target devices, wherein the same single base address is associated with each of the plurality of target devices.*” (Emphasis added).

As a preliminary matter, the Examiner has incorporated the rationale of the rejection of claim 1, but clearly only cited to Leung for the rejection of claims 16 and 25. As such, there is at least some ambiguity with respect to the rejection of claims 16 and 25

as to how Gutttag is being applied. Appellants request that the Examiner clarify his position on this rejection, as it appears to be incomplete. Otherwise Appellants respectfully request withdrawal of the rejection and allowance of the claim.

In rejecting claims 16 and 25, the Examiner again points to the memory modules of Leung as being target devices. However, as discussed in detail above, the memory modules cannot reasonably be considered the same as the target devices set forth in claim 16 and 25 without incurring inconsistency in the treatment of Leung. This is particularly the case with respect to claims 16 and 25 which recite an initiator device and target devices. The Examiner indicated that the I/O module 104 is an initiator in rejecting claims 16 and 25, but in the rejection of claims 17 and 26 indicated that the I/O module 104 is a target device. One of ordinary skill in the art would not consider the I/O module 104 as being both an initiator and a target device. Indeed, it is inconsistent and an unreasonable interpretation to do so. Accordingly, Appellants respectfully assert that the I/O module cannot reasonably be considered an initiator and the memory modules cannot reasonably be considered the target devices of claims 16 and 25. For at least this reason, Appellants respectfully request reversal of the rejection of claims 16 and 25.

Additionally, the Examiner has cited to first and second address fields for the memory modules in Leung as disclosing “a single base address associated with the plurality of target devices.” Even assuming *arguendo* that the memory modules could reasonably be considered target devices, the first and second field to which the Examiner refers are pertinent only to accessing a particular array or bank with a *single* memory module. *See, e.g.*, Leung, col 10, lines 21-25. Indeed, each of the memory modules is equipped with an independent address so that they can operate as independent units and each has its own, i.e., unique, base address. *See* Leung col. 4, lines 20-24. There is simply nothing in Leung that can reasonably be construed as disclosing a single base address associated with the plurality of target devices,” as set forth in claims 16 and 25.

Accordingly, Appellants respectfully assert that Leung taken alone or in combination with Guttag cannot support a *prima facie* case for obviousness under Section 103. Appellants respectfully request reversal of the rejection of claims 16 and 25, and all claims depending therefrom.

Dependent claims 17, 18, 26, and 27

Dependent claims 17 and 26 set forth wherein the target devices comprise input/output controllers. Dependent claims 18 and 27 set forth wherein the target devices comprise disk array controllers. As discussed above, the Examiner has inconsistently treated Leung and, thus, the terms of the claims. In particular, the Examiner insists that both memory modules and controllers of Leung can be deemed the equivalent of the target devices of the instant claims. Additionally, the Examiner indicates that the I/O module 104 of Leung can be the equivalent of both the target devices and the initiator. Appellants respectfully assert that such is not the case and, indeed, the Examiner's treatment only induces confusion. In particular, the I/O module and the memory modules cannot *both* reasonably be considered the same as the target devices of the claims, as they each perform distinctly different tasks and operations, and clearly are different devices. Moreover, even assuming *arguendo* that the Examiner's treatment was reasonable and accurate, Leung fails to disclose a plurality of I/O controllers and, as such, is still deficient in disclosing all the features of the claims. Accordingly, for at least these reasons, Appellants respectfully request reversal of the rejection of claims 17, 18, 26 and 27.

C. Ground of Rejection No. 3:

The Examiner rejected claim 24 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Carmichael. Specifically, the Examiner stated:

As per claim 24, Leung and Guttag disclose the claimed invention as detailed above in the previous paragraphs. However, Leung and Guttag do not explicitly

teach the communications bus comprises a Peripheral Component Interconnect (PCI) bus as recited in the claim.

Carmichael discloses the bus comprises a Peripheral Component Interconnect (PCI bus [*the bridge 36 may simply provide an extension of the processor's bus, or may buffer and extend the processor bus using an entirely different bus structure and protocol such as PCI*; col. 6, lines 46-50] to provide an extension of the processor's bus and to buffer and extend the processor (col. 6, lines 46-50).

Since the technology for implementing a data processing system using a PCI bus was well known in the art as evidenced by Carmichael, and since a PCI bus provides an extension of the processor's bus and to buffer and extend the processor, and artisan would have been motivated to implement a PCI bus in the data processing system of Leung and Gutttag. Thus it would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system fo Leung and Gutttag to include a PCI bus to provide an extension of the processor's bus and to buffer and extend the processor (col. 6, lines 46-50) as taught by Carmichael.

Office Action, pages 16-17.

Appellants respectfully assert that the Examiner has not established a *prima facie* case of obviousness with regard to dependent claim 24. Specifically, as described above, the Leung and Gutttag references clearly do not disclose those claim features attributed by the Examiner in the rejection of independent claim 16, upon which claim 24 depends. Carmichael fails to overcome the deficiencies of the Leung and Gutttag references with respect to claim 16. As such, the Leung, Gutttag, and Carmichael references, taken alone or in combination, do not disclose all the features of claim 16. As such, Appellants respectfully request withdrawal of the Section 103 rejection and allowance of claim 24 based on its dependency from claim 16.

D. **Ground of Rejection No. 4:**

The Examiner rejected claims 32 and 47 under 35 U.S.C. § 103(a) as being unpatentable over Gutttag and Gupta in view of Blaner.

Claim 32

With respect to claim 32, the Examiner stated:

As per claim 32, Gutttag discloses a plurality of devices having a common base address and a transaction request directed to the common base address *[each of a plurality of n processors has a predetermined plurality of corresponding memories, wherein the predetermined plurality of memories corresponding to each processor having a corresponding fixed base address; and a base address instruction executing on any one of the plurality of n processors generating the base address of the predetermined plurality of memories corresponding to the one processor; col. 172, ll 48-60; a plurality of corresponding memories having a unique addressable memory of a single memory space wherein these corresponding memories have a corresponding base address within the single memory address space; col. 5, lines 36-41]*. It is worth mentioning that Gutttag discloses a one to one correspondence between the “plurality of memories” and the “a corresponding fixed base address”.

Gupta discloses a computer comprising a memory [col. 1, lines 24-25]; a controller configured to logically divide the memory into a plurality of interleaved memory regions [Fig. 2]; and a plurality of devices, wherein each of the plurality of devices is associated with one of the interleaved memory regions and wherein each of the devices simultaneously accesses its associated interleaved memory region in response to a single transaction request [col. 6, lines 21-28; col. 16, ll 12-24] so that multiple CPUs tend not access the same memory bank at the same time (col. 6, lines 23-26)

Gupta further discloses each of the devices simultaneously accesses its associated interleaved memory region in response to a single transaction request *[two or*

more memory busses can each perform memory transactions simultaneously, with each memory bus coupled to one or more memory bus segments, wherein a single memory transaction can be active on a single memory bus segment at any given time, and each memory bus segment is coupled to one or more interleaved memory banks; col. 16, ll 15-23].

Blaner incontestably discloses each of the devices simultaneously accesses its associated interleaved memory region in response to a single transaction request [*interleaved memory banks wherein a group of the interleaved memory banks are accessible in parallel in response to a single access; col. 8, ll 22-26*] to allow simultaneous access to multiple pages of memory and reduce latency (col. 2, ll 65-67).

Since the technology for implementing a memory system with devices each simultaneously accessing its associated interleaved memory region in response to a single transaction requests was well known as evidenced by Blaner, an artisan would have been motivated to implement this feature in the system of Gupta to allow simultaneous access to multiple pages of memory and reduce latency. Thus, it would have been obvious to one of ordinary skill in the art at the time of invention by applicant to modify the system of Gupta in view of Blaner to include devices each simultaneously accessing its associated interleaved memory region in response to a single transaction requests because this would have allowed simultaneous access to multiple pages of memory and reduced latency (col. 2, ll 65-67) as taught by Blaner.

Office Action, pages 17-19 (emphasis in original).

Amended claim 32 recites, *inter alia*, “a plurality of devices having a *common base address*, wherein each of the plurality of devices is associated with one of the interleaved memory regions and wherein each of the devices *simultaneously accesses* its associated interleaved memory region *in response to a single transaction request directed to the common base address.*” (Emphasis added).

Appellants respectfully assert that Gupta fails to disclose all the elements of claim 32 and, further, that Blaner fails to obviate the deficiencies of Gupta. First, Gupta does not disclose a plurality of devices simultaneously accessing associated interleaved memory regions in *response to a single transaction*, contrary to the Examiner's assertions. As discussed above, Gupta is directed to a system wherein the memory is accessed simultaneously in response to *multiple transactions*. Indeed, the portions of Gupta relied on by the Examiner only disclose multiple simultaneous transactions resulting in simultaneous memory access with only a single transaction per bus at any one time. *See* Gupta, col. 6, lines 21-28; col. 16, lines 12-24. However, the Examiner has not cited to, and the Appellants are unaware of, anything in Gupta that can even remotely be considered as disclosing, teaching or suggesting a plurality of devices simultaneously accessing interleaved memory in response to *a single transaction*.

Additionally, there is nothing in Gupta with respect to the plurality of devices having a *common base address* and simultaneous access of the plurality of devices to associated interleaved memory region in response to *a single transaction request directed to the common base address*. Blaner does not overcome the deficiencies of Gupta in this respect.

Blaner is directed to the processing of storage keys in a computer that employs key-controlled storage protection. *See* Blaner, col. 1, lines 32-35. The two main features of the system in Blaner are 1) interleaved memory and 2) an associative multi-page key cache. *See id.* at col. 4, lines 60-64. The interleaved memory physically stores the keys in separate RAMs so that address A addresses the key for page p and the address A+1 addresses the key for page p+n. *See id.* at col. 4, line 64 through col. 5, line 1. According to Blaner, this organization allows for n keys of a plurality of pages to be fetched in one access. *See id.* at col. 5, lines 1-2. However, there is no mention in Blaner of a plurality of devices having a common base address wherein where *each*

simultaneously accesses its associated interleaved memory region in response to a single transaction request directed to the common base address, as set forth in claim 32. As such, Gupta and Blaner, taken alone or in combination, fail to disclose all the features of claim 32.

Accordingly, Appellants respectfully request withdrawal of the Section 103 rejection and allowance of claim 32. As such, Appellants respectfully request that the reversal of the Section 103 rejection based on Gupta and Blaner.

Claim 47

With respect to claim 47, the Examiner stated:

As per claim 47, Gutttag discloses a method of performing operations accessing memory comprising assigning a common base address to a plurality of disk drives to create a collective target group [col. 172, ll 48-60; col. 5, lines 36-41].

Gupta further discloses assigning a portion of data storage to each disk drive of the collective target group [Fig. 1] so that multiple CPUs tend not access the same memory bank at the same time (col. 6, lines 23-26].

Blaner discloses initiating an interleaved memory operation as a single request to the collective target group using the common base address, wherein each disk drive recognizes the common base address and simultaneously executed the memory operation on the portion of data storage to which the disk drive is assigned [*interleaved memory banks wherein a group of the interleaved memory banks are accessible in parallel in response to a single access*; col. 8, ll 22-26] to allow simultaneous access to multiple pages of memory and reduce latency (col. 2, ll 65-67).

Office Action, page 19 (emphasis in original).

Independent claim 47 recites, “*assigning a common base address to a plurality of disk drives to create a collective target group; assigning a portion of data storage to each disk drive...and initiating an interleaved memory operation as a single request to the collective target group using the common base address, wherein each disk drive recognizes the common base address and simultaneously executes the memory operation on the portion of data storage to which the disk drive is assigned.* (Emphasis added). In sharp contrast, the cited references, alone or in hypothetical combination do not disclose such features.

In particular, as discussed above, Guttag does not disclose assigning a common base address to a plurality of disk drives, contrary to the Examiner’s allegation. Rather, Guttag discloses unique addresses, i.e., different address for each of a plurality of memories. As stated in Guttag:

Each of the memories has a unique addressable memory portion of a single memory space. Each processor has a predetermined plurality of corresponding memories. These corresponding memories have a corresponding base address within said single memory address space.

Guttag, at col. 5, lines 36-41 (emphasis added). In rejecting claim 47, the Examiner referred to col. 172, lines 48-60 and col. 5, lines 36-41 of Guttag as disclosing the above-recited claim features. In pertinent part, the cited portion states:

a plurality of n processors, where n is less than m and each of said n processors has a predetermined plurality of corresponding memories, said predetermined plurality of memories corresponding to each processor having a corresponding fixed base address within said single memory address space, each of said processors capable of generating any address within said single memory address space for read/write access to data stored within said plurality of m memories.

Guttag, col. 172, lines 48-55 (emphasis added). As the above-quoted section clearly illustrates, Guttag discloses a plurality of processors that each have *their own* (i.e., *corresponding*) base address within *their own individual section* of the memory address space. *See also*, Guttag, col. 172, lines 43-47. That is, a memory space is divided into a plurality of *m* memories, each memory having a “*unique* addressable memory portion”. Consequently, the claim feature “assigning a common base address” is not taught by Guttag. Indeed, Guttag teaches the antithesis of this feature, in that each memory has its own, unique base address.

Gupta and Blaner do not overcome this deficiency of Guttag. In particular, Appellants are unaware of, and the Examiner has not cited to, any portion of Gupta and Blaner that overcome the deficiency of Guttag in this regard. As such, for at least this reason, Appellants assert that the Examiner has failed to present a *prima facie* case for obviousness under Section 103 and, as such request reversal of the rejection of claim 47.

Moreover, however, Gupta does not disclose the features attributed to it by the Examiner. In particular, Gupta does not disclose assigning a portion of data storage to each disk drive of the collective target group, as alleged. As set forth in the instant application, the data may be distributed among the collective target group having a common base address allowing individual target devices to store a portion of a large data request. *See, e.g.*, specification, page 9, ll. 18-32. This allows for multiple targets of the collective target group to simultaneously fulfill a single read or write request to reduce wait time. *Id.* In contrast, Gupta disclose different data hot spots being distributed across several memory banks to avoid having multiple CPUs trying to access the same memory at the same time. *See* Gupta, col. 6, line 23-29. Thus, Gupta simply distributes different sets of data in different memory banks to allow different requests to different data set to be performed at the same time. Therefore, Gupta cannot reasonably be considered as disclosing assigning a portion of data storage to each disk drive of the collective group, as set forth in claim 47.

In view of the foregoing, Guttag, Gupta, and Blaner, taken alone or in hypothetical combination fail to disclose all the features of claim 47 and cannot support a *prima facie* case of obviousness under Section 103. Accordingly, Appellants respectfully request reversal of the rejection of claim 47.

E. **Ground of Rejection No. 5:**

The Examiner rejected claim 48 under 35 U.S.C. § 103(a) as being unpatentable over Guttag and Gupta in view of Blaner and further in view of Olarig. Appellants respectfully traverse this rejection.

Claim 48 depends from independent claim 47. As set forth above, the Examiner has failed to present a *prima facie* case of obviousness with respect to claim 47. Appellants respectfully assert that claim 48 is allowable for at least the reasons set forth with respect to claim 47, base on the dependency from claim 47. Furthermore, Appellants assert that the Olarig reference fails to overcome the deficiencies of the Guttag, Gupta, and Blaner references with respect to claim 47. Accordingly, Appellants respectfully request reversal of the rejection of claim 48 based on its dependency from claim 47.

F. **Ground of Rejection No. 6:**

The Examiner rejected claims 33-38 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Gupta. Specifically, the Examiner stated:

As per claim 33, the rationale in the rejection of claim 1 is herein incorporated. However, Leung and Guttag do not explicitly teach a method comprising dividing a section of memory into a plurality of interleaved memory regions as required by the claims.

Gupta discloses a method comprising dividing a section of memory into a plurality of interleaved memory regions [col. 6, lines 20-24] so that multiple CPUs tend not to access the same memory bank at the same time (col. 6, ll 23-26).

Since the technology for implementing a data processing system with dividing a section of memory into a plurality of interleaved memory regions was well known as evidenced by Gupta, an artisan would have been motivated to implement this feature in the system of Leung and Guttag] so that multiple CPUs would not tend to access the same memory bank at the same time. Thus, it would have been obvious to one of ordinary skill in the art at the time of invention by Applicant to modify the system of Leung Guttag to include dividing a section of memory into a plurality of interleaved memory regions because this would have allowed multiple CPUs not to tend accessing the same memory bank at the same time (col. 6, ll 23-26) as taught by Gupta.

Office Action, pages 20-21. Appellants respectfully traverse this rejection.

Claim 33 recites, *inter alia*, “associating the plurality of target devices with a single base memory address, wherein the same single base memory address is associated with each of the plurality of target devices.” (Emphasis added). Claim 36 recites, *inter alia*, “code to configure the plurality of devices to associate a single base address with the plurality of devices, wherein the same single base address is associated with each of the plurality of devices.” (Emphasis added).

The Examiner admitted in the rejection of claim 1 that Leung fails to disclose a single base memory address being associated with each of the plurality of target devices. As set forth above, Guttag clearly does not to overcome the deficiencies of Leung in this regard. Specifically, Guttag discloses a system wherein processors have unique base addresses within memory space. *See* Guttag, col. 172, lines 43-55. This is the antithesis of the claimed feature having the same single base address being associated with each of

the plurality of devices. As such, the Leung and Guttag references, alone or in combination, do not disclose all the features of the independent claims 33 and 36.

Appellants respectfully assert that Gupta does not overcome this deficiency of the Leung and Guttag references. Gupta discloses an interleaved memory coupled with a plurality of memory buses over which *independent memory transactions* may simultaneously be performed, but where only a single memory transaction can be active on a memory bus at any given time. See Gupta, abstract; col. 6, lines 19-29; col. 16, lines 17-23. However, Gupta does not disclose that the same single base memory address is associated with each of a plurality of target devices, as set forth in claims 33 and 36. Accordingly, the Leung, Guttag and Gupta references, taken alone or in hypothetical combination, do not disclose all the features of independent claims 33 and 36.

As such, Appellants respectfully assert that a *prima facie* case of obviousness under Section 103 has not been established with regard to independent claims 33 and 36. Therefore, Appellants respectfully request reversal of the rejection under Section 103 and allowance of claims 33 and 36, as well as all claims depending therefrom.

G. **Ground of Rejection No. 7:**

The Examiner rejected claims 39-44 under 35 U.S.C. § 103(a) as being unpatentable over Leung in view of Guttag and further in view of Olarig. Appellants respectfully traverse this rejection.

As described above, the Leung and Guttag references clearly do not disclose those claim features attributed by the Examiner in the rejection of the independent claims 1, 16, and 25, upon which claims 39-44 depend, respectively. In view of the deficiencies of the Leung and Guttag references with respect to the independent claims, the Examiner's Section 103 rejections of claims 39-44, which are based upon the Examiner's mistaken

interpretation of these references, cannot establish a *prima facie* case of obviousness. As such, Appellants respectfully request reversal of the Section 103 rejections of claims 39-44.

H. **Ground of Rejection No. 8:**

The Examiner rejected claims 45 and 46 under 35 U.S.C. § 103(a) as being unpatentable over Gupta in view of Blaner and further in view of Olarig. Appellants respectfully traverse this rejection.

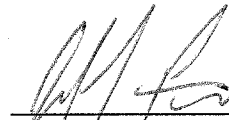
As stated above, the Gupta and Blaner references fail to disclose all the features of independent claim 32. The Olarig reference fails to overcome those deficiencies. As such, the Gupta, Blaner, and Olarig references, taken alone or in combination, fail to disclose all features of independent claim 32. In particular, Olarig does not disclose “a plurality of devices having a *common base address*, wherein each of the plurality of devices is associated with one of the interleaved memory regions and wherein each of the devices *simultaneously accesses* its associated interleaved memory region *in response to a single transaction request directed to the common base address.*” (Emphasis added). Appellants respectfully assert, therefore, that claims 45 and 46 are allowable based on their dependency from claim 32, and respectfully request reversal of the rejection.

Conclusion

Appellants respectfully submit that all pending claims are in condition for allowance. However, if the Examiner or Board wishes to resolve any other issues by way of a telephone conference, the Examiner or Board is kindly invited to contact the undersigned attorney at the telephone number indicated below.

Respectfully submitted,

Date: September 22, 2008



Jeffery R. Peterson
Reg. No. 62,488
FLETCHER YODER
P.O. Box 692289
Houston, TX 77269-2289
(281) 970-4545

Correspondence Address

IP Administration
Legal Department, M/S 35
HEWLETT-PACKARD COMPANY
P.O. Box 272400
Fort Collins, CO 80527-2400

8. **APPENDIX OF CLAIMS ON APPEAL**

Listing of Claims:

1. A method for transacting between an initiator device and a plurality of target devices, the method comprising:

 associating each of the plurality of target devices with a single base address,

 wherein the same single base address is associated with each of the plurality of target devices;

 sending a multicast transaction from the initiator device to the plurality of target devices, wherein sending the multicast transaction comprises sending a multicast transaction to the single base address associated with each of the plurality of target devices; and

 executing the transaction when the transaction is received by the plurality of target devices according to the configuration of the target devices.

2. The method of claim 1, wherein the associating each of the plurality of target devices comprises:

 assigning a base memory address to be shared by the plurality of target devices;

 and

 assigning a first portion of memory to a first target device of the plurality of target devices.

3. The method of claim 2, wherein the transaction is a read request for a block of stored data from memory, comprising:

recognizing the base memory address from the read request;

initiating a read operation by the plurality of target devices assigned to the base memory address;

fetching stored data from a portion of memory associated with each of the target devices, the data being concurrently fetched by each associated target device; and

sending the fetched data to the initiator device.

4. The method of claim 2, wherein the transaction is a write request for data to be stored in memory, comprising:

recognizing the base memory address from the write request;

initiating a write operation by the plurality of target devices assigned to the base memory address; and

writing data of the write request to a portion of memory associated with each target device, the data being concurrently written by each associated target device.

5. The method of claim 1, wherein the target devices comprise input/output controllers.

6. The method of claim 1, wherein the target devices comprise disk array controllers.

8. The method of claim 1, comprising a plurality of target groups.

9. A method for transacting data stored in memory between an initiator device and multiple target devices, the method comprising:

detecting a multicast transaction request;

accessing a first portion of memory by a first target device in response to the multicast transaction request; and

accessing a second portion of memory by a second target device concurrently with the access to the first portion of memory in response to the multicast transaction request, wherein the first and second portions of memory are accessed with a single base address associated with both the first target device and the second target device, wherein the first target device and the second target device are associated with the same single base address.

10. The method of claim 9, wherein the target devices comprise input/output controllers.

11. The method of claim 9, wherein the target devices comprise disk array controllers.

13. The method of claim 9, comprising accessing a plurality of target devices, wherein the plurality of target devices are divided into a plurality of groups, wherein each of the plurality of groups is associated with a single base memory address configured to address the target devices within that group.

16. A computer system comprising:

a bus;

an initiator device coupled to the bus, the initiator device configured to initiate a transaction request; and

a plurality of target devices coupled to the bus, wherein each of the plurality of target devices concurrently executes a portion of the transaction request, wherein the initiator device is configured to multicast the transaction request to the plurality of target devices using a single base address associated with the plurality of target devices, wherein the same single base address is associated with each of the plurality of target devices.

17. The computer system of claim 16, wherein the plurality of target devices comprise input/output controllers.

18. The computer system of claim 16, wherein the plurality of target devices comprise disk array controllers.

20. The computer system of claim 16, wherein the plurality of target devices comprise a target group.
21. The computer system of claim 20, comprising a plurality of target groups.
22. The method of claim 16, wherein the transaction is a multicast read request.
23. The method of claim 16, wherein the transaction is a multicast write request.
24. The computer system of claim 16, wherein the bus comprises a Peripheral Component Interconnect (PCI) bus.
25. A computer system comprising:
- a processor;
 - a bus coupled to the processor;
 - an initiator device coupled to the bus, the initiator device configured to issue a multicast transaction; and
 - a plurality of target devices coupled to the bus, the plurality of target devices configured to execute the multicast transaction with concurrent data responses from a plurality of interleaved memory regions, wherein the initiator device is configured to multicast the transaction request to the plurality of target devices using a single base address associated with the

plurality of target devices, wherein the same single base address is associated with each of the plurality of target devices.

26. The computer system of claim 25, wherein the target devices comprise input/output controllers.

27. The computer system of claim 25, wherein the target devices comprise disk array controllers.

29. The computer system of claim 25, wherein the plurality of target devices are divided into a plurality of target groups, wherein each of the target groups is associated with its own base address.

32. A computer comprising:

a memory;

a controller configured to logically divide the memory into a plurality of

interleaved memory regions; and

a plurality of devices having a common base address, wherein each of the

plurality of devices is associated with one of the interleaved memory

regions and wherein each of the devices simultaneously accesses its

associated interleaved memory region in response to a single transaction

request directed to the common base address.

33. A method comprising:

dividing a section of memory into a plurality of interleaved memory regions;

associating the plurality of interleaved memory regions with a plurality of target devices;

associating the plurality of target devices with a single base memory address,

wherein the same single base memory address is associated with each of the plurality of target devices; and

executing a memory access using the single base memory address.

34. The method of claim 33, wherein executing the memory access comprises executing a read operation.

35. The method of claim 33, wherein executing the memory access comprises executing a write operation.

36. A tangible machine readable medium comprising:

code to initialize a plurality of devices;

code to configure the plurality of devices to associate a single base address with

the plurality of devices, wherein the same single base address is associated with each of the plurality of devices; and

code to associate the single base address with a plurality of interleaved memory regions, wherein the same single base address is associated with each of the plurality of interleaved memory regions.

37. The tangible medium of claim 36, comprising:

code to issue a single read command comprising the single base address;

code to recognize the single base address as associated with the plurality of devices;

code to simultaneously execute a plurality of memory requests involving the plurality of devices;

code to receive data from the plurality of devices; and

code to write the received data to a bus.

38. The tangible medium of claim 36, comprising:

code to issue a write command comprising the single base address;

code to recognize the base address as associated with the plurality of devices; and

code to simultaneously write to the plurality of devices.

39. The method of claim 1, wherein sending the multicast transaction comprises sending the multicast transaction from a disk drive controller to a plurality of disk drives.

40. The method of claim 1, wherein sending the multicast transaction comprises sending the multicast transaction from a SCSI controller to a plurality of SCSI devices.

41. The computer system of claim 16, wherein the initiator device comprises a SCSI controller.

42. The computer system of claim 16, wherein the plurality of target devices comprises a plurality of disk drives.

43. The computer system of claim 25, wherein the initiator device comprises a SCSI controller.

44. The computer system of claim 25, wherein the plurality of target devices comprises a plurality of SCSI hard drives.

45. The computer of claim 32, wherein the controller comprises a SCSI controller.

46. The computer of claim 32, wherein the plurality of devices comprises a plurality of SCSI hard drives.

47. A method for performing operations accessing memory comprising:

assigning a common base address to a plurality of disk drives to create a
collective target group;
assigning a portion of data storage to each disk drive of the collective target
group; and
initiating an interleaved memory operation as a single request to the collective
target group using the common base address, wherein each disk drive
recognizes the common base address and simultaneously executes the
memory operation on the portion of data storage to which the disk drive is
assigned.

48. The method of claim 47, wherein initiating an interleaved memory operation
as a single request comprises initiating a multicast transaction from a SCSI controller
over a multicast bus.

9. **EVIDENCE APPENDIX**

None.

10. **RELATED PROCEEDINGS APPENDIX**

None.